

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Robert M. Ellis Confirmation No.: 8266
Serial No.: 10/713,563 Examiner: Michael C. Krofcheck
Filed: November 13, 2003 Group Art Unit: 2186
For: METHOD AND APPARATUS FOR MAINTAINING DATA
DENSITY FOR DERIVED CLOCKING
Date: February 5, 2008

**AMENDED APPEAL BRIEF
FILED IN THE USPTO VIA EFS**

AMENDED APPELLANT'S BRIEF

UNDER 37 CFR §41.37

Appeal is taken from the Examiner's Final Office Action mailed April 3, 2008 and Advisory Action mailed June 12, 2008, finally rejecting claims 1-15, 25 and 26 in the instant application.

This Amended Appeal Brief is in furtherance of the Notice of Appeal filed in this case on July 3, 2008 and Notification of Non-Compliant Appeal Brief mailed January 26, 2009.

The fees required under §41.37(a)(2) and any required petition for extension of time for filing this Brief and fees therefor are dealt with in the accompanying TRANSMITTAL OF AMENDED APPEAL BRIEF.

This Brief contains these items under the following headings, and in the order set forth below.

TABLE OF CONTENTS

- I. REAL PARTY IN INTEREST
- II. RELATED APPEALS AND INTERFERENCES
- III. STATUS OF CLAIMS
- IV. STATUS OF AMENDMENTS
- V. SUMMARY OF CLAIMED SUBJECT MATTER
- VI. GROUNDS OF REJECTION TO BE REVIEWED UPON APPEAL
- VII. ARGUMENT
- VIII. CLAIMS APPENDIX
- IX. EVIDENCE APPENDIX
- X. RELATED PROCEEDINGS APPENDIX

I. REAL PARTY IN INTEREST

37 CFR §41.37(c)(1)(i)

Robert M. Ellis and Intel Corporation are the real parties in interest.

II. RELATED APPEALS AND INTERFERENCES

37 CFR §41.37(c)(1)(ii)

The Board's decision in the present Appeal will not directly affect, or be directly affected, or have any bearing on any other appeals or interferences known to the appellant, or to the appellant's legal representative.

III. STATUS OF CLAIMS

37 CFR §41.37(c)(1)(iii)

Status of All the Claims:

1. Claims presented: 1-15, 25 and 26.
2. Claims withdrawn from consideration but not cancelled: 16, 17 and 22-24.
3. Claims canceled: 18-21 and 27-28
4. Claims pending: 1-15, 25 and 26.

of which:

- a. claims allowed: None
- b. claims rejected: 1-15, 25 and 26.

All the rejected claims, namely 1-15, 25 and 26 are being appealed. The appealed claims are eligible for appeal, having been finally rejected.

IV. STATUS OF AMENDMENTS

37 CFR §41.37(c)(1)(iv)

An amendment was filed on June 3, 2008 in response to the Final Office Action dated April 3, 2008. An Advisory Action dated June 12, 2008 indicated that the amendment of June 3, 2008 was entered for the purposes of appeal. There have been no amendments filed subsequent to the Advisory Action dated June 12, 2008.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

37 CFR §41.37(c)(1)(v)

The claimed subject matter relates to memory systems, components, and methods and, more particularly, to a method and apparatus for maintaining data density for a derived clocking technology.

Claim 1 recites a method comprising calculating an achieved data transition density for at least one data lane in a point-to-point memory channel having a plurality of data lanes, the achieved data transition density calculated over greater than two clock cycles (see, e.g., specification, p. 3, l. 23 – p. 4, l. 26); and transmitting a synchronization signal on the at least

one data lane responsive to the achieved transition density (see, e.g., specification, p. 3, ll. 8-22, and p. 5, ll. 5-16).

Claim 3 recites storing a desired data transition density for the at least one data lane (see, e.g., specification, p. 4, ll. 1-6); and comparing the achieved data transition density to the desired data transition density (see, e.g., specification, p. 4, ll. 1-6).

Claim 4 recites that transmitting a synchronization signal on the at least one data lane responsive to the achieved transition density comprises transmitting a synchronization signal on all the data lanes if the achieved data transition density is less than the desired data transition density on the at least one data lane (see, e.g., specification, p. 4, ll. 16-24).

Claim 5 recites a memory channel comprising a host and a plurality of DIMMs connected in a point-to-point fashion, wherein the host includes a processor (see, e.g., specification, p. 2, ll. 14-17); an outbound data channel and an inbound data channel, each having a plurality of data lanes (see, e.g., specification, p. 3, ll. 23-29); at least one transition detection circuit configured to detect whether an achieved data transition density on at least one data lane is less than a desired data transition density for the at least one data lane (see, e.g., specification, p. 3, l. 23- p. 5, l. 22); and a transition generator configured to transmit a synchronization signal on the at least one data lane if the achieved transition density is less than the desired data transition density (see, e.g., specification, pp. 5-8); wherein the achieved transition density is measured over greater than two clock cycles (see, e.g., specification, p. 4, ll. 1-6).

Claim 8 recites that the at least one transition detection circuit comprises a plurality of data transition detectors, each configured to detect a data transition on a corresponding data lane (see, e.g., specification, p. 3, l. 23 – p. 6, l. 34); a clock cycle counter (see, e.g., specification, p. 4, ll. 7-13); a plurality of data transition counters, each configured to count the data transitions detected by a corresponding data transition detector, and configured to be reset by the clock cycle counter (see, e.g., specification, p. 4, ll. 1-6); a logic block configured to signal when at least one of the plurality of data transition counters counts, during a time period defined by the clock cycle counter, a number of data transitions detected by the corresponding data transition

detector that is less than the desired data transition density (see, e.g., specification, p. 4, ll. 16-27).

Claim 14 recites an article of machine-readable code, embodied on a machine-readable medium, that when executed causes a machine to perform processes comprising storing a desired data transition number (see, e.g., specification, p. 4, ll. 1-6); storing a clock cycle number that is greater than two (see, e.g., specification, p. 4, ll. 7-14); for a data lane in a point-to-point memory channel, recording a measured data transition number over a period of clock cycles equal to the clock cycle number (see, e.g., specification, p. 4, ll. 1-6); comparing the measured data transition number to the desired data transition number (see, e.g., specification, p. 4, ll. 1-6); and transmitting a synchronization signal on the data lane if the measured data transition number is less than the desired data transition number (see, e.g., specification, p. 5, ll. 5-16).

Claim 25 recites that transmitting the synchronization signal on the at least one data lane comprises transmitting the synchronization signal having a number of transitions to cause the achieved data transition density for the at least one data lane to be greater than or equal to a desired data transition density (see, e.g., specification, p. 3, ll. 8-15, and p. 4, ll. 16-27).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

37 CFR §41.37(c)(1)(vi)

Claims 1-7 and 25-26 are rejected under 35 USC 103(a) as being unpatentable over U.S. Pub. No. 2002/0112119 to Halbert et al. (“Halbert”) and U.S. Pub. No. 2005/0005046 to Bashirullah et al. (“Bashirullah”).

Claims 8, 9, 11, and 12 are rejected under 35 USC 103(a) as being unpatentable over Halbert and Bashirullah as applied to claim 5, and further in view of U.S. Pat. No. 5,533,696 to Boggs, et al. (“Boggs”).

Claims 10 and 13 are rejected under 35 USC 103(a) as being unpatentable over Halbert, Bashirullah and Boggs as applied to claim 8, and further in view of applicant’s admitted prior art (“AAPA”).

Claims 14-15 are rejected under 35 USC 103(a) as being unpatentable over Halbert, Bashirullah, Boggs and U.S. Pub. No. 2004/0056782 to Bliss et al. (“Bliss”).

VII. ARGUMENT

37 CFR §41.37(c)(1)(vii)

I. Grouping of Claims

For purposes of the rejections under 35 U.S.C. § 103(a), the claims include five groups of claims. Claims 1 and 2 are grouped together. Claim 3 and 4 are grouped together. Claims 5-7 and 26 are grouped together. Claims 8, 9, 11, and 12 are grouped together. Claim 25 is argued separately. Claims 10 and 13 are grouped together. Claims 14 and 15 are grouped together.

II. Rejections under 35 USC 103(a) over Halbert and Bashirullah

A. Halbert and Bashirullah in General

Halbert focuses on a dual port buffer to memory interface. In particular, Halbert describes interfaces between a memory controller 200 and memory modules 150 through a point-to-point bus.¹

Bashirullah focuses on an adaptive bandwidth bus. In particular, Bashirullah describes a bus that can switch from a current mode to a voltage mode of operation in response to bus signal activity.² Bashirullah describes a FIFO 410 for a data bit Din[0] at the input of a bus 100. Data passes through the FIFO 410 before passing on to the repeaters 440, 450, etc. The bits of the FIFO 410 are combined to determine if there are any transitions within the FIFO 410.

If there are any transitions, the control signal generator 430 asserts a control signal on the control line 105. In response, the repeaters 440, 450 switch to a current mode from a voltage mode to enter a higher bandwidth state. If there are no transitions, the control signal

¹ Halbert, Title, ¶[0054], and FIGS. 10-11.

² Bashirullah, Abstract and ¶[0046].

causes the repeaters to enter the lower bandwidth voltage mode. As a result, when transitions are present in the data, the bus 100 operates in the higher bandwidth current mode. When transitions are not present, the repeaters operate in the lower bandwidth and lower power voltage mode. Accordingly, power consumption is reduced.³

With respect to the combination of Halbert and Bashirullah, the Examiner notes that Halbert does not teach the transition densities, synchronization signals, or the like related to transitions.⁴ Halbert is apparently included to teach a point-to-point memory channel.⁵ Accordingly, although the arguments below refer to Bashirullah with respect to transitions, the Applicant asserts that neither Halbert, nor the knowledge of one skilled in the art teaches or suggests such elements as well.

B. Claims 1 and 2

Insofar as claims 1 and 2 have been rejected under the same grounds, Applicant argues the claims as a group. The arguments below are made with respect to claim 1 on behalf of the group.

Claim 1 recites:

A method comprising:

calculating an achieved data transition density for at least one data lane in a point-to-point memory channel having a plurality of data lanes, the achieved data transition density calculated over greater than two clock cycles; and

transmitting a synchronization signal on the at least one data lane responsive to the achieved transition density.

1. The Examiner is interpreting “synchronization signal” of claim 1 in isolation of the specification

The Examiner has two alternative interpretations of a synchronization signal. First,

³ See Bashirullah, FIGS. 4 and 5, and ¶[0056]-[0062].

⁴ See, for example, Final Office Action dated April 3, 2008 (“Final Office Action”), p. 3.

⁵ See, for example, Final Office Action p. 3.

the Examiner argues that the “synchronization signal is the data bits outputted from the repeaters on the data lane.”⁶ Second, the Examiner argues that the “control signals [are] the synchronization signals.”⁷ However, both of these interpretations have been made in isolation of the specification. In particular, both interpretations have nothing to do with synchronization.

Referring to the specification, for example:

“In order to maintain the phase relationship between the data eye and the derived clock, a certain minimum number of data transitions in a fixed time period are required. In other words, a minimum transition density must be maintained. The minimum transition density is typically on the order of 5 data transitions for every 128 clock cycles, but it may be more or less depending on system requirements. To guarantee that the minimum number of data transitions occurs, the host 210 may periodically send a synchronization data stream that contains the required data transitions on the outbound data path 215b.”⁸

“DoSyncOB is asserted when one of the outbound data paths requires synchronization signals after the programmed number of clock cycles and DoSyncIB is asserted when one of the inbound data paths requires synchronization signals after the programmed number of clock cycles.”⁹

That is, regardless of its particular form, the synchronization signal aids in synchronization. In the examples above, the synchronization signal guarantees that the minimum number of transitions occurs so that the phase relationship between the data eye and the derived clock can be maintained.

In contrast, the interpretations of the Examiner eliminate any meaning of the word “synchronization.” For example, using the first interpretation of the data bits as the synchronization signal, the interpretation is unreasonably broad. The Examiner is apparently arguing that *any* data bit transmitted over the bus 100 of Bashirullah is a synchronization signal. However, this would mean that all of the data transmitted over the bus is the synchronization signal.

What is missing in Bashirullah is the association of those data bits with some form of

⁶ Final Office Action, p. 11.

⁷ Advisory Action dated June 12, 2008 (“Advisory Action”), p. 3.

⁸ Specification, p. 3, ll. 3-10.

⁹ Specification, p. 5, ll. 7-10.

synchronization. The Applicant is not arguing that data bits cannot be a synchronization signal; rather that there must be some indication that bits that are interpreted as a synchronization signal have some relationship to synchronization, however broadly interpreted, that is consistent with the specification.

Turning to the alternative control signal interpretation of the synchronization signal, the control signal 105 of Bashirullah, again, has nothing to do with synchronization. As described above, the purpose of the control signal is only to control the bandwidth of the bus 100. It does not affect any synchronization.

For example, comparing the bus upon which Bashirullah improves to the bus of Bashirullah, there is no control signal as the synchronization signal. Data bits are propagated since the repeaters of the bus remain in the current mode. The addition of the control signal in Bashirullah allows the repeaters to enter the lower power voltage mode. The control signal does not affect the data bits in any way different from the conventional bus described in Bashirullah. That is, the improvement of Bashirullah only reduces power consumption.

If there was some synchronization aspect of the control signal of Bashirullah, there should be some effect on some synchronization if it is not present. However, there is no difference in the performance of the conventional bus and the bus of Bashirullah with the exception of reduced power. The transmitted data remains the same. Accordingly, the control signal of Bashirullah has no effect on synchronization.

2. *The data bits of Bashirullah cannot be the “synchronization signal” of claim 1*

Claim 1 recites “transmitting a synchronization signal on the at least one data lane responsive to the achieved transition density.” That is, the transmission of the synchronization signal is in response to the achieved transition density. Assuming for the sake of argument that the data bits are the synchronization signal, the transmission of the data bits must be in response to the achieved transition density.

In contrast, in Bashirullah, the actual transmitting of the data bits is not affected by the transition density. Consider the two cases of the current mode and the voltage mode of Bashirullah, described above. The data bits are being transmitted continuously, regardless of

the transition density. For example, if the data includes transitions, the repeaters are switched to the current mode so that the transitions are propagated. If the data does not include transitions over a length of C_p+1 of the FIFO 410¹⁰, that means that the data bits within the FIFO are either all 1's or all 0's. The repeaters are switched to the voltage mode; however, the 1's or 0's are still transmitted.

Thus, regardless of the transition density, the data bits are transmitted. Accordingly, the data bits cannot be the synchronization signal since whether the data bits are transmitted is independent of and not responsive to the transition density.

As the rejection of claim 1 relies on this unreasonably broad interpretation of synchronization signal, and at least part of that interpretation conflicts with the language of claim 1, the Applicant respectfully requests that the Board reverse the rejection of claims 1 and 2.

C. Claims 3 and 4

Insofar as claims 3 and 4 have been rejected under the same grounds, Applicant argues the claims as a group. The arguments below are made with respect to claim 3 on behalf of the group. However, as claim 3 is dependent on claim 1, the arguments above with respect to the deficiencies of the rejection of claim 1 also apply to claim 3.

Claim 3 recites “storing a desired data transition density for the at least one data lane.” In Bashirullah, no desired data transition density is stored. The Examiner argues that “as the threshold is set [at] 1, it must be kept track of somewhere and compared to the occurred transitions.”¹¹

¹⁰ See Bashirullah, ¶[0060]-[0062].

¹¹ See Final Office Action, p. 4.

1. No transition density is “desired”

In Bashirullah, there is not a desired transition density. The Examiner points to 1 transition as the desired amount; however, there is no preference for 0, 1, or more than 1 transition. That is, Bashirullah dutifully transmits the data regardless of the transition density. The transition density of “1” is only used to control the bandwidth of the repeaters. It has no bearing on whether transitions are desired or not.

In contrast, consider the example described above in the specification where a certain transition density is desired to maintain a phase relationship between a data eye and a clock. In this example, a particular transition density is desired, namely, a transition density above which the phase relationship of the data eye and the clock are maintained.

Although the term “desired” is not limited to the above example, there must be some reason why a particular transition density is preferred over another in order for a transition density to be a desired transition density. In Bashirullah, there is no reason why 0, 1, or more than 1 transition is desired over some other transition density.

2. There is no “storing” of a transition density

In Bashirullah, the transition density of 1 is not stored. Rather, it is inherent in the logic that combines the bits of the FIFO to determine if there are any transitions. Thus, in direct contrast to the Examiner’s assertion that the transition density must be “kept track of somewhere,” nothing keeps track of the desired transition density.

Moreover, it should be noted that the transition density in Bashirullah *must* be 1. For example, if it is 0, the bus 100 always operates in contrast to the desired goal of Bashirullah. If the transition density is 2 or more, then the data is distorted. For example, if a long sequence of 0’s is followed by a long sequence of 1’s, as the transition propagates through the FIFO, the repeaters will not enter the current mode as there is only 1 transition. Accordingly, the single transition will be distorted due to the lower bandwidth of the voltage mode, rendering Bashirullah unfit for its intended purpose. Since Bashirullah requires a transition density of 1, there is no reason to store that value as it will never change.

Accordingly, there is no desired transition density in Bashirullah. The Applicant respectfully requests that the Board reverse the rejection of claim 3.

D. Claims 5-7 and 26

Insofar as claims 5-7, and 26 have been rejected under the same grounds, Applicant argues the claims as a group. The arguments below are made with respect to claim 5 on behalf of the group.

Claim 5 recites:

A memory channel comprising:

a host and a plurality of DIMMs connected in a point-to-point fashion, wherein the host includes a processor;

an outbound data channel and an inbound data channel, each having a plurality of data lanes;

at least one transition detection circuit configured to detect whether an achieved data transition density on at least one data lane is less than a desired data transition density for the at least one data lane; and

a transition generator configured to transmit a synchronization signal on the at least one data lane if the achieved transition density is less than the desired data transition density;

wherein the achieved transition density is measured over greater than two clock cycles.

1. No transition density is “desired”

As described above with respect to claim 3, there is no desired data transition density. Bashirullah has no preference for a particular transition density. Since there is no desired data transition density, a synchronization signal cannot be transmitted by the transition generator depending on whether the achieved transition density is less than a non-existent desired data transition density.

2. No synchronization signal is transmitted when the achieved transition density is less than the desired data transition density

The Examiner argues that the “zero transition density of Bashirullah being below the threshold density of one causes the control signal (synchronization signal) to switch the operation to [the] voltage mode.”¹²

As described above with respect to claim 1, the Examiner made particular interpretations of a synchronization signal, namely, that the synchronization signal was the control signal. In this instance, the Examiner is apparently arguing that the synchronization signal is the state of the control signal of Bashirullah that causes the repeaters to enter the voltage mode.

Similar to the discussion of the synchronization signal with respect to claim 1, there is no indication that a signal that switches the repeaters of Bashirullah into the voltage mode has anything to do with synchronization. For example, in order for the control signal of Bashirullah to be generated to switch the repeaters to the voltage mode, a long sequence of 1's or 0's must fill the FIFO entirely. Since the change in the control signal follows after the last transition in the FIFO, the state of the control signal to switch the repeaters to the voltage mode has no effect on those transitions propagating through the repeaters.

Since the FIFO is filled with 1's or 0's, there are no transitions to introduce a time relationship into the data. Since there is no time relationship, there cannot be any synchronization. That is, no synchronization occurs with respect to a sequence of 1's or a sequence of 0's since there are no transitions within the sequence that can be affected by the control signal.

Accordingly, Halbert and Bashirullah fail to teach a desired data transition density or a synchronization signal being transmitted when the achieved data transition density is less than the desired data transition density. The Applicant respectfully requests that the Board reverse the rejection of claims 5-7 and 26.

¹² Advisory Action, p. 5.

E. Claim 25

As claim 25 is dependent on claim 1, the arguments above with respect to the deficiencies of the rejection of claim 1 also apply to claim 25. Moreover, claim 25 recites that “transmitting the synchronization signal on the at least one data lane comprises transmitting the synchronization signal having a number of transitions to cause the achieved data transition density for the at least one data lane to be greater than or equal to a desired data transition density.” That is, the achieved data transition density becomes equal to or exceeds the desired data transition density *because of* the transmission of the synchronization signal. In other words, before the transmission of the synchronization signal, the achieved data transition density was less than the desired data transition density.

The Examiner argued that the detection of a change in the data bits in Bashirullah causes the control signal to be asserted, which causes the repeaters to change to the current mode so that the transitions are propagated, which causes the transition density to go from zero to one.¹³ However, the Examiner is apparently referring to two different data transition densities. In the rejection of claim 1, the Examiner apparently referred to the achieved data transition density as what was determined using the FIFO.¹⁴ However, the transition density which the Examiner states transitions from zero to one is apparently a transition density that is beyond the FIFO, somewhere near the repeaters. In other words, the Examiner is citing a *different* transition density and is not citing the transition density where the achieved transition density was calculated.

Moreover, in order for the control signal to be sent to propagate the transitions over the repeaters, the achieved data transition density must have already been calculated as greater than 1. Thus, the achieved data transition density does not become greater than or equal to the desired transition density because of the synchronization signal. In contrast, the causality is reversed. The cited synchronization signal of Bashirullah is transmitted *because* the achieved data transition density is greater than or equal to one.

¹³ Final Office Action, p. 11-12.

Accordingly, in addition to the deficiencies of Halbert and Bashirullah with respect to claim 1, Halbert and Bashirullah does not teach that the synchronization signal causes the achieved data transition density to be greater than or equal to the desired data transition density. The Applicant respectfully requests that the Board reverse the rejection of claim 25.

III. Rejections under 35 USC 103(a) over Halbert, Bashirullah, and Boggs

A. Boggs in General

Boggs introduces a transmission speed detection circuit.¹⁵ Over a particular time period, transitions of the data are counted. Based on how many transitions there are during the time period, a determination of a speed of the data can be made.¹⁶ Transitions are counted until the time expires or the threshold of transitions is attained. In particular, the output of the asynchronous counter 52, 54 is only asserted *after* a sufficient number of transitions are counted to exceed the threshold.¹⁷

B. Claims 8, 9, 11, and 12

Insofar as claims 8, 9, 11, and 12 have been rejected under the same grounds, Applicant argues the claims as a group. However, as claims 8, 9, 11, and 12 are dependent on claim 5, the arguments above with respect to the deficiencies of the rejection of claim 5 also apply to claims 8, 9, 11, and 12. The arguments below are made with respect to claim 8 on behalf of the group.

Claim 8 recites, *inter alia*:

a clock cycle counter; and

a plurality of data transition counters, each configured to count the data transitions detected by a corresponding data transition detector, and configured to be reset by the clock cycle counter.

¹⁴ Final Office Action, p. 3.

¹⁵ Boggs, Abstract.

¹⁶ See Boggs, col. 6, l. 48 - col. 7., l. 27.

¹⁷ Boggs, col. 6, ll. 54-67.

1. The Examiner's articulated reasoning to combine Boggs with Halbert and Bashirullah is contradicted by Bashirullah

The Examiner argued that the counters of Boggs would be combined with Halbert and Bashirullah with the motivation “to provide increased flexibility in the driver in the combination of Halbert and Bashirullah, which also allows for monitoring of transitions on a *longer* scale for better overall efficiency.”¹⁸

However, the exact opposite would occur. In fact, if the time period over which transitions are monitored is increased, there is an increased likelihood that a transition would be detected. As described above, once a transition is detected in Bashirullah, the repeaters transition to the higher power current mode. The longer the time period, the more likely the repeaters are in the higher power mode, and the less likely the repeaters remain in the lower power voltage mode. As a result, increasing the time period for counting transitions *increases* power consumption thereby decreasing overall efficiency in direct contrast to the Examiner's cited motivation.

2. The Examiner's articulated reasoning to combine Boggs with Halbert and Bashirullah disregards the teachings of Bashirullah and Boggs

The Examiner's motivation disregards the transition detection techniques of Bashirullah and Boggs. For example, in Bashirullah, the rolling window for detection implemented by the FIFO is necessary for proper operation. The depth of Cp +1 of the FIFO allows the repeaters time to transition from the voltage mode to the current mode.¹⁹ Without this delay, the first few transitions would be distorted during the time that the repeaters are transitioning to the current mode. Even if the length of time over which transitions are detected in Bashirullah is increase, a FIFO would still be required to buffer the data after that transition is detected.

¹⁸ Final Office Action, p. 7, emphasis added.

¹⁹ Bashirullah, ¶[0060].

With respect to Boggs, two outputs indicate whether a particular number of transitions have been detected. The output of the counter 52, 54 indicates whether a number of transitions have been detected within the current time period.²⁰ However, the counter is reset at the beginning of each time period.²¹ Thus, at the beginning of each time period, this output would switch the repeaters of Bashirullah into the voltage mode, regardless of whether there are transitions.

Boggs does teach a latched output from latch 56. However, this output does not change until the beginning of a subsequent time period. As described above, in Bashirullah, the transition to the current mode occurs before a bit is output to the repeaters. If the output of the latch 56 of Boggs is used, there may not be enough time to allow the repeaters to transition to the current mode before that bit is output to the repeaters. For example, consider a bit arriving at the beginning of a time period. The output of the latch 56 does not transition to the end of the time period, during which the bit is propagating through the required FIFO. As a result, the transition may be output just as the latch 56 indicates that there are transitions, giving the repeaters no time to transition to the current mode, thereby distorting the transition. Accordingly, Bashirullah would become unfit for its intended purpose of conserving power without affecting the data.

Furthermore, even if the circuitry of Boggs were added to Bashirullah, this again is in contrast to the goal of Bashirullah of reducing power consumption. As described above, the FIFO of Bashirullah must remain. Accordingly, the circuitry of Boggs would be in addition to that of Bashirullah. The additional circuitry, in particular the additional circuitry that must operate at the higher speed to detect the incoming transitions, must be added to each bus line Din[0..N] of Bashirullah. Accordingly, a vast amount of circuitry would be unnecessarily added to Bashirullah, thereby increasing the power consumption in the combination.

²⁰ Boggs, col. 6, ll. 48-53.

²¹ Boggs, col. 6, ll. 42-47.

Accordingly, one skilled in the art would not add the counters 52, 54 of Boggs to Halbert and Bashirullah as argued by the Examiner. The Applicant respectfully requests that the Board reverse the rejection of claims 8, 9, 11, and 12.

IV. Rejections under 35 USC 103(a) over Halbert, Bashirullah, Boggs, and the AAPA

Insofar as claims 10 and 13 have been rejected under the same grounds, Applicant argues the claims as a group. However, as claims 10 and 13 are dependent on claims 8 and 11, respectively, the arguments above with respect to the deficiencies of the rejection of claim 5 also apply to claims 10 and 13.

Claims 10 and 13 were rejected with the addition of the AAPA to Halbert, Bashirullah, and Boggs. However, the addition of the AAPA does not overcome the deficiencies of Halbert, Bashirullah, and Boggs with respect to the rejection of parent claims 8 and 11, respectively, described above. In particular, the Examiner cited the AAPA to teach AND gates and NAND gates.²² However, the addition of AND or NAND gates as described in the AAPA does not teach a desired data transition density, nor the transmission of a synchronization signal as described above with respect to parent claim 5, nor does the addition cure the conflict between the Examiner's articulated reasoning and the teaching of Halbert, Bashirullah, and Boggs. The Applicant respectfully requests that the Board reverse the rejection of claims 10 and 13.

V. Rejections under 35 USC 103(a) over Halbert, Bashirullah, Boggs, and Bliss

Insofar as claims 14 and 15 have been rejected under the same grounds, Applicant argues the claims as a group. Claims 14 and 15 were rejected with the addition of Bliss; however Bliss was apparently introduced only for the purpose of citing machine-readable code. The addition of Bliss does not cure the deficiencies of Halbert, Bashirullah, and Boggs as described above.

In particular, claim 14 recites, *inter alia*:

²² Final Office Action, p. 7.

storing a desired data transition number;
storing a clock cycle number that is greater than two;
for a data lane in a point-to-point memory channel, recording a measured data transition number over a period of clock cycles equal to the clock cycle number;
comparing the measured data transition number to the desired data transition number;
and
transmitting a synchronization signal on the data lane if the measured data transition number is less than the desired data transition number.

As described above with respect to claim 1, a synchronization signal, given its broadest reasonable interpretation consistent with the specification, is not transmitted. As described above with respect to claim 3, there is no desired data transition density, nor is such desired density stored. Similarly, in claim 14, there is no desired data transition number, nor is such a number stored. As described above with respect to claim 5, no synchronization signal is transmitted when the achieved transition density is less than the desired data transition density. Similarly, in claim 14, no synchronization signal is transmitted if the measured data transition number is less than the desired data transition number.

Claims 1, 3, and 5 were rejected over Halbert and Bashirullah. However, as described above with respect to claim 8, the addition of Boggs did not cure the deficiencies of claim 5 and similar elements of claim 1 and 3. The addition of Bliss was apparently to teach a computer readable medium.²³ Accordingly, Halbert, Bashirullah, Boggs, and Bliss fail to teach each and every element of claim 14. The Applicant respectfully requests that the Board reverse the rejection of claims 14 and 15.

²³ Final Office Action, p. 10.

CONCLUSION

For the foregoing reasons, Applicant requests that the Board reverse the Examiner's rejections of Applicant's claims.

Respectfully submitted,

MARGER JOHNSON & McCOLLOM, P.C.

A handwritten signature in black ink, appearing to read 'Derek Meeker', is written over a horizontal dotted line.

Derek Meeker
Reg. No. 53,313

MARGER JOHNSON & McCOLLOM, P.C.
210 SW Morrison Street, Suite 400
Portland, OR 97204
503-222-3613

VIII. CLAIMS APPENDIX

37 CFR §41.37(c)(1)(viii)

1. (Previously presented) A method comprising:

calculating an achieved data transition density for at least one data lane in a point-to-point memory channel having a plurality of data lanes, the achieved data transition density calculated over greater than two clock cycles; and

transmitting a synchronization signal on the at least one data lane responsive to the achieved transition density.
2. (Previously presented) The method of claim 1, wherein calculating an achieved data transition density for the at least one data lane comprises:

counting how many times a data transition occurs on the at least one data lane during a predetermined number of clock cycles, the predetermined number of clock cycles being greater than two.
3. (Original) The method of claim 2, further comprising:

storing a desired data transition density for the at least one data lane; and

comparing the achieved data transition density to the desired data transition density.
4. (Original) The method of claim 3, wherein transmitting a synchronization signal on the at least one data lane responsive to the achieved transition density comprises:

transmitting a synchronization signal on all the data lanes if the achieved data transition density is less than the desired data transition density on the at least one data lane.

5. (Previously Presented) A memory channel comprising:

a host and a plurality of DIMMs connected in a point-to-point fashion, wherein the host includes a processor;

an outbound data channel and an inbound data channel, each having a plurality of data lanes;

at least one transition detection circuit configured to detect whether an achieved data transition density on at least one data lane is less than a desired data transition density for the at least one data lane; and

a transition generator configured to transmit a synchronization signal on the at least one data lane if the achieved transition density is less than the desired data transition density;

wherein the achieved transition density is measured over greater than two clock cycles.

6. (Original) The memory channel of claim 5, wherein the at least one transition detection circuit is located on the host.

7. (Original) The memory channel of claim 5, wherein the at least one transition detection circuit is located on a corresponding one of the plurality of DIMMs.

8. (Original) The memory channel of claim 5, wherein the at least one transition detection

circuit comprises:

a plurality of data transition detectors, each configured to detect a data transition on a corresponding data lane;

a clock cycle counter;

a plurality of data transition counters, each configured to count the data transitions detected by a corresponding data transition detector, and configured to be reset by the clock cycle counter;

a logic block configured to signal when at least one of the plurality of data transition counters counts, during a time period defined by the clock cycle counter, a number of data transitions detected by the corresponding data transition detector that is less than the desired data transition density.

9. (Original) The memory channel of claim 8, wherein the clock cycle counter and the plurality of data transition counters are programmable.

10. (Original) The memory channel of claim 8, wherein the logic block comprises an AND gate and a plurality of NAND gates.

11. (Original) The memory channel of claim 5, wherein the at least one transition detection circuit comprises:

a plurality of data transition detectors, each configured to detect a data transition on a corresponding data lane;

a clock cycle counter;

a plurality of data transition counters, each configured to count the data transitions detected by a corresponding data transition detector, and configured to be reset by the clock cycle counter;

a first logic block configured to signal when at least one of the plurality of data transition counters corresponding to the data lanes on the outbound data path counts, during a time period defined by the clock cycle counter, a number of data transitions detected by the corresponding data transition detector that is less than the desired data transition density; and

a second logic block configured to signal when at least one of the plurality of data transition counters corresponding to the data lanes on the inbound data path counts, during a time period defined by the clock cycle counter, a number of data transitions detected by the corresponding data transition detector that is less than the desired data transition density.

12. (Original) The memory channel of claim 11, wherein the clock cycle counter and the plurality of data transition counters are programmable.

13. (Original) The memory channel of claim 11, wherein the first logic block and the second logic block comprise a plurality of NAND gates and an AND gate.

14. (Previously presented) An article of machine-readable code, embodied on a machine-readable medium, that when executed causes a machine to perform processes comprising:

storing a desired data transition number;

storing a clock cycle number that is greater than two;
for a data lane in a point-to-point memory channel, recording a measured data transition number over a period of clock cycles equal to the clock cycle number;
comparing the measured data transition number to the desired data transition number;
and
transmitting a synchronization signal on the data lane if the measured data transition number is less than the desired data transition number.

15. (Previously presented) The article of machine-readable code of claim 14, that when executed, causes the machine to perform processes further comprising:

transmitting a synchronization signal on the data lane if the measured data transition number is less than the desired data transition number.

16. (Withdrawn) A method comprising:

operating a selected data lane from a point-to-point memory channel having a plurality of data lanes in an inverted mode according to a preselected data inversion scheme.

17. (Withdrawn) The method of claim 16, wherein operating the selected data lane in the inverted mode comprises:

applying data inversions simultaneously to both a receiver and a transmitter of the selected data lane in a node of the point-to-point memory channel.

18. - 21. (Cancelled)

22. (Withdrawn) A machine-readable medium, that when read, causes a machine to perform processes comprising:

operating a selected data lane of a point-to-point memory channel in an inverted mode according to a preselected data inversion scheme.

23. (Withdrawn) The machine-readable medium of claim 22, wherein operating the selected data lane of the point-to-point memory channel in an inverted mode according to a preselected data inversion scheme comprises:

applying data inversions simultaneously to a plurality of receivers and a plurality of transmitters that correspond to the selected data lane, wherein the data inversions are applied according to the preselected data inversion scheme.

24. (Withdrawn) The machine-readable medium of claim 22, that when read, causes the machine to perform processes further comprising:

loading a bit sequence that represents the preselected data inversion scheme into a wraparound shift register.

25. (Previously presented) The method of claim 1, wherein transmitting the synchronization signal on the at least one data lane comprises:

transmitting the synchronization signal having a number of transitions to cause the

achieved data transition density for the at least one data lane to be greater than or equal to a desired data transition density.

26. (Previously presented) The memory channel of claim 5, wherein for each of the at least one transition detection circuit:

the transition detection circuit is configured to detect whether an achieved data transition density on a corresponding one of the at least one data lane is less than the desired data transition density for the corresponding data lane.

27.-28. (Canceled)

IX. EVIDENCE APPENDIX
37 CFR §41.37(c)(1)(ix)

NONE

X. RELATED PROCEEDINGS APPENDIX
37 CFR §41.37(c)(1)(x)

NONE